




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA




South African National Grid

Application Porting and Support

21.9.2011

Bruce Becker | bbecker@csir.co.za
Coordinator – SAGrid | <http://www.sagrid.ac.za>





Outline

- Reminder of SAGrid infrastructure
- Challenges identified by SAGrid
- SAGrid users and user Communities
- Applications : the good news and the bad
- Strategic considerations for the region's future







NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT





UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA

 Functional Sites
  Core Services
  Deployment In progress
  Proposed extensions
  External Infrastructure


 Existing SANReN backbone
 Construction in progress
 Proposed extensions
 TENET/SEACOM 10Gbps

7 sites providing
~1300 CPU cores
~15 TB perm. storage
17 in SAGridOps

 EUMed
 EGI
 EELA
 CERN



7(+3) sites
~1300 CPU cores
~20 TB perm. Storage
~17 permanent staff



SAGrid and User Communities


- SAGrid is a *shared, federated* infrastructure
- Resources provided by *institutes*, not users -
 - SAGrid sites responsible for *operations* only
- Some "self-supported" communities :
 - Provide their own resources :
 - ATLAS VO (2 sites)
 - ALICE VO (2 sites)
 - E-NMR (2 sites)
- These consist of ~50 % of the usage of the grid, but constitute only few individuals.




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT




UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



The long tail... that wags the dog





Users and User Communities

- SAGrid consists of compute resources at the universities – which host the ***single*** users
- Not organised into VO's - SAGrid has a catch-all VO to contain them and their applications
- User support and identification is done by first-contact at the site, or by dedicated training events (EPIKH)
- In March 2011 after internal discussion, some challenges were identified, specifically related to users




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

UCT CERN



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



Challenge 3 : Inclusion

- Many universities are fully integrated into the grid, but we are missing some
- Next EPIKH school to be held at Stellevbosch – aim to encourage uptake and usage of the services
- Still missing : "the coast" ; "FET"s
- Message to users must be clear : no matter whether your particular institute ***provides resources*** to the grid, ***you can use it.***




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

UCT CERN



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



Challenge 4 : Diversity


- SAGrid needs to be able to cover all use cases , measured by many metrics
 - **Compute aspects** (IPC, paralellism, complexity)
 - **Data aspects** (scale, complexity, distribution, ownership)
 - **Application aspects** (diversity, support, interaction)
 - **Interoperability** – all services need to be able to talk to all other services, with no bottlenecks
- This is a distinguishing factor of the grid – and should be held up as a **design** benefit



NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



Challenge 5 : Barriers to Entry

- Interaction with and maintenance of the grid can be made far more intuitive and easy
 - Infrastructure-level policy on IaaS
 - Virtual Research Environments for user communities
- Basic grid web portal at <http://ui.sagrid.ac.za>
- But ! community-specific web-based VRE's and application-specific Science Gateways needed




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

UCT CERN




UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



The good news



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



The good news - inclusion


- EPIKH Africa 5 school :
 - Great effort to develop coherent deployment strategy, streamlined development
 - obtained > 20 requests for applications
 - Ported 9 during the school – 6 already in an AppDB
 - 5 different countries represented, apart from SA
 - Most of the work done in **remote**, using variation of EPIKH porting procedure
 - Almost all applications deployed with documentation, sample JDL and script
 - <https://ops.sagrid.ac.za/trac/repo/demo/>



NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



The good news - diversity


- Wide diversity in applications -
 - Applications in many fields
 - Chemistry, CFD, Mathematics, biology, climate, e-learning, earth/livestock observation, astrophysics
 - Real-time applications using SAGrid infrastructure
 - Trad. batch applications with high data throughput
 - Trad. HPC applications
 - Single, short-run applications, with constant user interaction
- All saw the beauty of the grid :)



NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



The good news - sustainability

- A lot was learned and implemented from previous experience
 - Application identification and porting (EPIKH/GILDA)
 - MPI integration (I2G/SEEGrid)
 - User grid execution environment (PLGrid)
 - User support (AfricaROC)
 - Documentation/publication (EPIKH/CHAIN AppDB)
- Applications from EUMed, EGI AppDB were ported "as-is" where possible




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

UCT CERN



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



Some particular examples - **CORDEX**

- COordinated Regional climate Downscaling Experiment

<http://www.meteo.unican.es/en/projects/CORDEX>

- A ready-made VRC for climate, using WRF4
- University of Cape Town, a coordinating partner – Climate Systems Analysis Group CSAG :
<http://www.csag.uct.ac.za/>




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

UCT CERIN



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



SAGrid and CSAG

- A good example of a community bringing its own resources
- Very good internal support for the application
- Proposal :
 - deploy the supported middleware on their hardware, SAGrid provides operational support
 - CSAG provides shared compute resources, participates to application support for their community
- The African core of a climate VRC for CHAIN to develop




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

UCT CERN



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



Some particular examples : single users

- The grid "caught" some users at the participating institutes, using already ported applications
 - OpenFOAM, Gaussian, R, MrBayes, GATE, AutoDock, HTK...
- Pros -
 - Users are immediately running jobs, and justifying the investment
 - Users are providing good feedback on how to improve the services and the infrastructure.




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

UCT CERN




UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



The bad news





The bad news

- User expectations
 - Infrastructure is not tuned to users' needs
 - Simple storage and computing **access** is hindered by complex user **interface**
 - Users not sure how to proceed once they have the basics
- Manpower involved is limited – need to prioritise applications
- Support is focussed almost exclusively in the porting phase – once application is ported, no support SLA




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

UCT CERN



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



Plans to address these issues

- Definition of a *User Support activity* in SAGrid
- Assignment of *priority* to applications, feedback to users
- Online *dashboard* of porting progress
- Clearer dissemination of what users can *expect* from the infrastructure and user support after the application porting
- Sites to define more rigorous, graduated *SLA's* for usage.




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

UCT CERN



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



Plans to address user inclusion


- Almost all users were disappointed by the UI, but enthusiastic about science gateways.
- We are starting to see a stabilisation in the number of applications we need to support
- Some users need refinement on their usage of the infrastructure-level services:
 - Workflow engines
 - Metadata catalogues
 - Distributed data storage
- Bi-monthly feedback session to identify and fix issues



NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



Will it scale ?

- The porting and support activity cannot be exclusively undertaken by SAGrid → current model is entirely unsustainable, even in SA.
- But capacity is out there -
 - Many applications are self-supported, even developed through other institutes... AIMS, universities, ICTP, etc
- The only way to sustain user support is to open it up to those who have the skills and will
- The social web is a key part of this...




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

UCT CERIN



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



Thank you

- Special thanks to SAGrid Operations Team
 - Albert van Eck, Tiaan Bezuidenhout (UFS)
 - Timothy Carr, Andrew Lewis (UCT)
 - Adrian Snyman, Francois Wolmarans, Stavros Lambropoulos (UJ)
 - Sean Murray (iThemba LABS)
 - Fourie Joubert (UP)
 - Thabo Molambo, Hannes Kriel (NWU)
 - Scott Hazelhurst, Sahal Yacoob (Wits)
- ... and EPIKH – Riccardo Rotondo, Jorge Sevilla




NORTH-WEST UNIVERSITY
YUNIBESITHI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

UCT CERIN



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA



Thank you

<http://www.sagrid.ac.za/>

<http://www.facebook.com/SAGrid>

<http://www.twitter.com/TheSAGrid>

Bruce Becker for
SAGrid Operations Team
SAGrid Joint Research Unit
bbecker@csir.co.za

